

LIST OF INVENTORS' NAMES AND ADDRESSES

Kentaro SHIMADA, Tokyo, JAPAN;

Akiyoshi HASHIMOTO, Kawasaki, JAPAN.

United States Patent Application

Title of the Invention

**STORAGE HAVING LOGICAL PARTITIONING
CAPABILITY AND SYSTEMS WHICH INCLUDE
THE STORAGE**

Inventors

Kentaro SHIMADA,

Akiyoshi HASHIMOTO.



13281
021904
U.S.PTO

STORAGE HAVING LOGICAL PARTITIONING CAPABILITY AND SYSTEMS WHICH INCLUDE THE STORAGE

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention relates to a storage which is connected to a network and used, in particular to a NAS.

Description of the Related Art

As a method of improving performance of an information processing system, the number of computers provided in the information processing system is simply increased. However, if a large number of computers are used, considerable time and labor are required for supervision of the respective computers, and a total area for installing the computers and total power consumed by the computers increase considerably. In order to solve this problem, there has been devised a technique for setting a high performance computer, logically partitioning resources of the computer into plural sections, and using the respective partitioned sections as a virtual computer independently. This is called logical partitioning (LPAR) of a computer. For example, an example of such a logical partitioning technique is disclosed in JP-A-2003-157177 (corresponding US Patent Publication No. 2003/0097393).

By virtualizing one computer look as if it is functioning as a large number of computers according to the logical

partitioning, a flexible operation becomes possible. For example, various operating systems can be used on the respective virtual computers freely and boot up and shutdown or failure management can be performed independently for each virtual computer. In addition, the number of physical machines is small, which is advantageous in terms of system management, an installation area for the machines, power consumption thereof, and the like. However, in the LPAR in the conventional computer, although resources such as a processor and a memory in the computer are logically partitioned and allocated to the respective virtual computers, concerning a storage connected to the computer, a storage area provided in the storage is simply partitioned and the partitioned storage areas are allocated to the virtual computers, respectively. Nothing is further taken into account specifically.

On the other hand, as a form of using a storage, other than a form in which the storage is directly connected to one computer (hereinafter referred to as "host" in some cases) and used, there is a form in which the storage is shared by plural computers via a network. In the latter form, in particular, a storage which has an interface as a form of a file system, that is, to which file access is possible from the computers, is called a network attached storage (hereinafter referred to as "NAS").

Data is exchanged between the NAS and the hosts by a form

of file having a name and a structure which are recognized by an operating system running on the host. Therefore, in addition to a disk drive for storing data and a control unit therefor, the NAS has a processor and a memory, which are used for translating file input/output to and from the host into data input/output to and from the disk drive.

SUMMARY OF THE INVENTION

The NAS is originally based upon an idea that it is more advantageous to provide a storage shared over a network than to set a storage individually for plural hosts. Thus, the NAS itself has to perform control for sharing a storage area or the like among the plural hosts. In addition, when a certain host reads and writes a large quantity of data, almost the entire processing ability of the NAS is consumed, and ability for inputting and outputting data to and from the other hosts declines. Moreover, when data destruction or failure of the NAS occurs due to operation mistake or the like of a certain host, the data destruction or failure may affect data used by the other hosts.

It is an object of the present invention to reduce control for sharing when the NAS is shared by plural hosts and eliminate mutual interference among the hosts so as to guarantee ability of input/output and localize data destruction or failure. In addition, it is another object of the present invention to

improve a usability of resources such as a processor and a memory in the NAS.

In order to attain the above-mentioned objects, logical partitioning is performed in the NAS. More specifically, there is provided a storage which is connected to a network and receives file access, in which resources held by the storage, for example, disk drives, interfaces with the network, processors controlling file access, and the like are logically partitioned by a control unit provided in the storage to enable respective logical partitions (virtual storages) to operate independently.

Further, it is also possible that the control unit has plural processors and these processors divides logical partitioning or performs logical partitioning as a whole.

Moreover, the control unit may have a supervising terminal to input information necessary for logical partitioning from this supervising terminal. In this case, it is also possible that only an access characteristic of a computer using the storage are inputted to the supervising terminal, and the supervising terminal calculates information necessary for logical partitioning from the access characteristic and communicates the information to the storage.

Moreover, a computer using the storage may also function as a supervising terminal.

BRIEF DESCRIPTION OF THE DRAWINGS

In the accompanying drawings:

Fig. 1 is a diagram showing an example of a structure of a NAS;

Fig. 2 is a diagram showing an example of allocation of a file control memory and a cache memory;

Fig. 3 is a diagram showing an example of allocation of a file control processor and a disk control processor;

Fig. 4 is a diagram showing an example of allocation of an inter-processor communication unit;

Fig. 5 is a diagram showing an example of allocation of groups of disk drives;

Fig. 6 is a diagram showing a concept of user authentication of a host system according to logical partitioning of a host interface unit and the file control processor;

Fig. 7 is a diagram showing an example of a structure of a NAS;

Fig. 8 is a diagram showing an example of a setting input screen for logical partitioning of the NAS; and

Fig. 9 is a diagram showing an example of logical partitioning information of respective resources of the NAS.

28

DESCRIPTION OF THE PREFERRED EMBODIMENT

Embodiments of the present invention will be hereinafter described with reference to the accompanying drawings. Note

that it is needless to mention that the present invention is not limited to descriptions of the embodiments described below.

Fig. 1 is a diagram showing an example of an embodiment of a NAS to which the present invention is applied. The NAS includes: two host interface units 100 which are connected to a host (hereinafter also referred to as "host system"); three file control processors 101 which translate an input/output request by a unit of file from the host into an input/output request of data by a unit of block; two file control memories 102 which store information necessary for translating an input/output request by a form of file into an input/output request of data by a unit of block (hereinafter also referred to as "translation control data"); four groups of disk drives 107 serving as storage media; three disk control processors 104 which control input/output of data to and from the groups of disk drives 107; two cache memories 106 which temporarily store data inputted to or outputted from the groups of disk drives 107; two disk interface units 105 which connect the disk control processors 104 and the groups of disk drives 107; and two inter-processor communication units 103 which connect the disk control processors 104 and the file control processors 101.

Here, it is also possible that plural disk drives are included in the groups of disk drives 107 and each group takes a RAID configuration. In addition, the "block" is a

predetermined unit which is used when the disk control processors 104 store data in disk drives. In general, 512B is adopted as the block. In addition, the "translation control data" is information indicating a correspondence relation between a file name used in a file system and a position of a file from its top, and the block. In general, the translation control data is expressed as a link structure such as I-Node or a data structure such as an address translation table.

Note that, in this specification, identical devices are denoted by identical reference numerals, and alphabets such as "a" and "b" are attached to the identical reference numerals when the identical devices are distinguished. In addition, the above-mentioned numbers of the respective devices are only examples and do not limit the present invention.

In Fig. 1, an input/output request by a unit of file, which is sent from the host system to the NAS, is processed in the NAS as described below.

First of all, a request for starting file reference (open) designating a file name is sent to the NAS from the host. Next, an actual input/output request for data is sent, and finally, a request for ending file reference (close) is sent.

These requests are received by any one of the host interface units 100 and transferred to any one of the file control processors 101. The file control processor 101 checks the file name requested by the host with reference to the translation

control data stored in the file control memories 102, records the file name started to be used, and translates a data input/output request for the file into a data input/output request to the groups of disk drives 107 in which the data is stored.

The translated data input/output request is sent to any one of the disk control processors 104 via any one of the inter-processor communication units 103. In addition, in the case in which necessary translation control data is not stored in the file control memories 102, the file control processor 101 requests translation control data stored in a predetermined storage area of the groups of disk drives 107 from any one of the disk control processors 104 via the inter-processor communication unit 103.

With respect to the data input/output request (including a request for translation control data) received from the any one of the file control processors 101 via the inter-processor communication unit 103, the disk control processor 104 checks if the data is stored in any one of the cache memories 105. When the data is stored in any one of the cache memories 105, the disk control processor 104 applies writing or reading of the requested data to the cache memory 105.

Thereafter, in the case of writing, the disk control processor 104 returns a result to the effect that writing is completed, or in the case of reading, returns a result to the

effect that reading is completed together with the read-out data to the file control processor 101, which sent the input/output request, via the inter-processor communication unit 103. The file control processor 101 processes the returned result and data and sends the result (data, a report on completion of processing, etc.) to the host system, which sent the input/output request to the NAS, via the host interface unit 100 which received the input/output request.

On the other hand, in the case in which the data requested by the file control processor 101 is not stored in all the cache memories 105, the disk control processor 104 specifies in which portion of the groups of disk drives 107 the requested data is stored, reads out the data from the portion of the groups of disk drives 107 via one of the disk interface units 106, and stores the data in one of the cache memories 105.

Thereafter, the disk control processor 104 applies reading or writing of the requested data to the cache memory 105 in which the data is stored. The subsequent processing is the same as the processing described above.

Note that the data stored in the cache memory 105 is written back to the groups of disk drives 107, for example, when a fixed time has elapsed or when a free space of the cache memory 105 becomes insufficient.

In this embodiment, for example, the above-mentioned kinds of processing are performed independently from each other

according to a unit (logical partition) partitioned by a logical partition boundary as shown in Fig. 1. When the host interface units 100, the file control processors 101, the file control memories 102, the inter-processor communication units 103, the disk control processors 104, the cache memories 105, the disk interface units 106, and the groups of disk drives 107, which are physical resources for processing allocated to the respective logical partitions, are allocated to each logical partition once, the devices are used solely for processing of the logical partition. More specifically, a file control processor 101a and a disk control processor 104c, which are allocated to different logical partitions in Fig. 1, do not exchange the input/output request of data as described above.

In addition, resources (e.g., file control memory 102b), which are shown across the logical partition boundaries in Fig. 1, are used for each logical partition boundary with a capacity or the like thereof logically partitioned at a rate allocated in advance. In this way, the logical partitions operate as virtual NASs independent from each other.

The processing for partitioning and allocating the physical resources to the respective logical partitions is actually executed by the file control processors 101 and the disk control processors 104. As a method of controlling logical partitioning, two methods described below are conceivable.

In a first method, the file control processors 101 and

the disk control processors 104 divide the control for the logical partitioning of several physical resources and control the logical partitioning in association with each other as a whole.

For example, the file control processors 101 perform processing for allocating the host interface units 100, the file control processors 101, and the file control memories 102. This processing is hereinafter referred to as "file control hypervisor".

In addition, the disk control processors 104 perform processing for allocating the inter-processor communication units 103, the disk control processors 104, the cache memories 105, the disk interface units 106, and the groups of disk drives 107. This processing is hereinafter referred to as "disk control hypervisor". The file control hypervisor executed by the file control processors 101 and the disk control hypervisor executed by the disk control processors 104 cooperate with each other to perform the each allocation processing. Details of cooperation will be described later. Note that the file control hypervisor may be executed by any one of the file control processors 101, for example, the file control processor 101a or may be executed by plural file control processors 101, for example, the file control processors 101a and 101b. This is also true for the disk control hypervisor.

In a second method, the two kinds of control processors

cooperate to control logical partitioning of all the physical resources. More specifically, all the file control processors 101a to 101c and the disk control processors 104a to 104c perform processing for allocating logical partitions of all the resources of the NAS (integrated hypervisor).

More specifically, for example, hypervisor operating on each processor realizes the logical partitioning as described below.

First, the hypervisor makes resources other than IO processing resources in a logical partition, to which each processor is allocated, invisible for basic IO processing software (BIOS) running on the processor. For example, in Fig. 1, the file control processor 101a is physically connected to the host interface unit 100b. However, in the case in which the logical partition is set to be partitioned by an illustrated dotted line, the hypervisor makes the host interface unit 100b invisible.

More specifically, in the case in which a privileged instruction for checking resources connected to the processor and available resources is executed in a BIOS, the hypervisor generates interrupt in terms of software according to execution of the privileged instruction to shift the execution to the hypervisor. The hypervisor checks resources allocated to a logical partition to which the processor belongs, sets a result of the privileged instruction such that only resources allocated

to the logical partition are visible, and returns to the BIOS in which the interrupt was generated.

In this way, each processor handles only resources of a logical partition to which the processor belongs, and separation of logical partitions is realized.

There are two kinds of memories, the host interface units 100 and the inter-processor communication units 103, which are provided with plural communication channels, among the resources. In such a case, it is sufficient to control an amount of resources to be shown to processors in the respective logical partitions (in the case of the memory, capacities of the memory in a start physical address and an end physical address, and in the case of the communication channel, the number of channels indicated by a set of physical numbers of channels).

In addition, concerning the processor itself, in the case in which each processor is allocated to one logical partition completely, it is sufficient to occupy the processor with the logical partition.

On the other hand, it is also conceivable to allocate one certain processor to two or more logical partitions and cause the logical partitions to share the processor by determining their sharing ratios. In such a case, it is conceivable to implement timer interrupt in each processor in terms of hardware to make arrangement such that the hypervisor is started up at each fixed time by the timer interrupt.

The hypervisor started up by the timer interrupt measures to which extent processing of a logical partition has been performed by the processor, determines a logical partition to be processed next in accordance with the predetermined sharing ratios, and shifts the execution of the processor to processing of the logical partition. In this way, one processor can be partitioned at predetermined ratios and allocated to two or more logical partitions.

Note that, as a method of realizing the hypervisor, other than the above-mentioned example, for example, it is also possible to provide special purpose hardware which is connected to the respective processors and performs resource supervision, a special purpose co-processor which is controlled by a small-sized micro program, and the like to realize the control of logical partitioning.

In addition, information on logical partitioning, for example, information designating a processor, a memory, a communication unit, and the like used in a logical partition 1, is stored in any one or more of the file control memories 102, the cache memories 105, the disk drives of the groups of disk drives 107, or the other storages. Each kind of hypervisor reads out the information to thereby apply designation of logical partitioning to the BIOS or the like. Note that this information is set via a supervising terminal to be described later.

Concerning the hypervisor realized as described above,

in the case in which the file control hypervisor is operated by the file control processor and the disk control hypervisor is operated by the disk control processor, the file control hypervisor performs the processing for allocating the host interface units 100, the file control processors 101, and the file control memories 102, the disk control hypervisor performs the processing for allocating the inter-processor communication units 103, the disk control processors 104, the cache memories 105, the disk interface units 106, and the groups of disk drives 107, and the two kinds of hypervisor are associated with each other.

More specifically, in designating logical partitioning with a supervising terminal to be described later, designation applied to logical partitioning with respect to the file control hypervisor and designation applied to logical partitioning with respect to the disk control hypervisor are performed in association with each other. Alternatively, depending upon a manner of designation with respect to logical partitioning, as described later, the file control hypervisor and the disk control hypervisor are adapted to automatically make adjustment each other such that a logical partition conforms to a request for the designated logical partitioning.

In the case of integrated hypervisor, for example, respective processors starting up the integrated hypervisor share information on allocation of all resources to respective

logical partitions, the hypervisor determines resources to be used by a processor, which has started up the hypervisor, with reference to the allocation information and executes allocation processing.

Next, a specific example of allocation of the resources of the NAS to logical partitions will be described. The allocation of the resources will be hereinafter described with the case in which the NAS shown in Fig. 1 is logically partitioned into two logical partitions (logical partition 1 and logical partition 2) as an example. However, any number of logical partitions will do. In addition, in the following description, the file control hypervisor and the disk control hypervisor are described as performing logical partitioning in association with each other. However, logical partitioning may be performed by the integrated hypervisor. Moreover, as a matter of expression, if hypervisor performs logical partitioning, processing by the hypervisor is actually executed by a processor which performs processing of each hypervisor.

Fig. 2 is a diagram showing an example of allocation of the file control memory 102 and the cache memory 105 to logical partitions. For example, when a request of a host system using the logical partition 1 attaches importance to a read data transfer rate, the disk control hypervisor increases an amount of allocation of the cache memory 105 to the logical partition 1 such that requested data is stored in the cache memory 105

as much as possible. In this case, an amount of allocation of the file control memory 102 to the logical partition 1 may be small.

In association with the above, the file control hypervisor decreases the amount of allocation of the file control memory 102 to the logical partition 1 to allocate a larger storage capacity to the logical partition 2. Consequently, it is possible to improve utilization of the file control memory 102 in the NAS as a whole.

On the other hand, if a request of the host system using the logical partition 1 attaches importance to a response time, the file control hypervisor allocates a larger storage capacity of the file control memory 102 to the logical partition 1 such that translation control data is stored in the file control memory 102 as much as possible. In this case, the amount of the cache memory 105 allocated to the logical partition 1 may be small. Consequently, the disk control hypervisor can allocate a larger capacity of the cache memory 105 to the logical partition 2, and it is possible to improve utilization of the cache memory 105 in the NAS as a whole.

In addition, in the case in which an input/output request of the host system using the logical partition 1 mainly concerns random access to data scattered in a large area of the groups of disk drives 107 provided in the NAS, it is difficult to store all pieces of information on the access scattered in a large

area even if large capacities of the file control memory 102 and the cache memory 105 is allocated to the logical partition 1. Thus, an effect of the allocation is small. Therefore, in such a case, the file control hypervisor and the disk control hypervisor reduce the amounts of allocation of the file control memory 102 and the cache memory 105 to the logical partition 1 and allocate a large storage capacity to the logical partition 2 which is the other logical partition to thereby improve a utilization of the cache memory 105 and the like.

Conversely, in the case in which an input/output request of the host system using the logical section 1 mainly concerns sequential access to data stored in continuous ranges of the groups of disk drives 107 provided in the NAS, it is possible to specify in advance information necessary for access and data to be read in advance in the NAS itself. Therefore, it is conceivable that the file control hypervisor and the disk control hypervisor perform allocation processing to increase amounts of allocation of the file control memory 102 and the cache memory 105 to the logical partition 1 such that the information and the data can be stored in the file control memory 102 and the cache memory 105 sufficiently.

Fig. 3 is a diagram showing an example of allocation of the file control processor 101 and the disk control processor 104 to logic partitions.

In the case in which an input/output request from a host

system using the logical partition 1 is a request for a small number of large file accesses, an amount of processing for translating file input/output executed by the file control processor 101 into data input/output is not so large. Therefore, an amount of allocation of the file control processor 101 to the logical partition 1 may be small.

In this case, the file control hypervisor reduces the amount of allocation of the file control processor 101 to the logical partition 1 (more specifically, an occupation ratio of the processor) and allocates relatively a larger amount of processor resources than that of the logical partition 1 to the logical partition 2 which is the other logical partition. Consequently, it becomes possible to improve a utilization of the file control processor 101 provided in the NAS.

In addition, in this case, since an amount of data of a file is large, an amount of processing of data input/output executed by the disk control processor 104 allocated to the logical partition 1 increases. Therefore, the disk control hypervisor increases an amount of allocation of the disk control processor 104 to the logical partition 1.

Further, in the case in which an input/output request of the host system using the logical partition 1 is a request for a large number of small file accesses, an amount of processing for translating file input/output executed by the file control processor 101 into data input/output increases. Thus, the file

control hypervisor increases an amount of allocation of the file control processor 101 to the logical partition 1.

In this case, since an amount of data of a file is small in itself, an amount of processing for data input/output executed by the disk control processor 104 allocated to the logical partition 1 is not so large. Thus, the disk control hypervisor reduces the amount of allocation of the disk control processor 104 to the logical partition 1 and increases an amount of allocation of the disk control processor 104 to the logical partition 2. Consequently, it becomes possible to improve a utilization of the disk control processor 104 in the NAS.

Moreover, in the case in which the host system using the logical partition 1 does not require a high performance NAS, the file control hypervisor and the disk control hypervisor perform allocation processing so as to reduce the amounts of allocation of the file control processor 101 and the disk control processor 104 to the logical partition 1. Conversely, in the case in which the host system using the logical partition 1 requires a high performance NAS, the file control hypervisor and the disk control hypervisor perform allocation processing so as to increase the amounts of allocation of the file control processor 101 and the disk control processor 104 to the logical partition 1.

Fig. 4 is a diagram showing an example of allocation of the inter-processor communication unit 103 to logical

partitions. In the case in which an input/output request from a host system using the logical partition 1 is a request for a large file sequential access, the disk control hypervisor performs allocation processing so as to increase an amount of allocation of the inter-processor communication unit 103 to the logical partition 1 (more specifically, a communication bandwidth) and preserve data communication ability between the file control processor 101 and the disk control processor 104 (in other words, data communication ability between the host interface unit 100 to the cache memory 105).

In addition, if an input/output request of the host system using the logical partition 1 is a request for a small file sequential access, the amount of allocation of the inter-processor communication unit 103 to the logical partition 1 may not be large. Moreover, if a request from the host system is a request for a random access, performance of a virtual NAS of the logical partition 1 from the viewpoint of the host system is not significantly affected by the amount of allocation of the inter-processor communication unit 103 to the logical partition 1. Therefore, in these cases, the disk control hypervisor performs allocation processing so as to reduce the amount of allocation of the inter-processor communication unit 103 to the logical partition 1 and increase allocation thereof to the other logical partition (here, the logical partition 2) and improve utilization of the inter-processor communication

unit 103 in the NAS.

Fig. 5 is a diagram showing an example of allocation of the groups of disk drives 107 to logical partitions. If a host system using the logical partition 1 prefers large storage capacity, the disk control hypervisor allocates the groups of disk drives 701, which has a RAID 5 configuration with a high storage capacity efficiency (in Fig. 7, while the number of disk drives in which data is stored is three, the number of disk drives in which parity is stored is one, and a capacity efficiency is 75%), to the logical partition 1. In this case, a speed of rotation of a disk drive may not be so fast, for example, 7,500 rpm.

On the other hand, if the host system using the logical partition 1 prefers good access performance, the disk control hypervisor allocates the group of disk drives 701, which has a RAID1 configuration allowing accessibility to be improved (in Fig. 7, since identical data is duplicated and stored in two disk drives, a storage capacity efficiency is 50%, but since both the two disk drives can be used for the identical data, total accessibility is twice as high as that of one disk drive), to the logical partition 1. Note that, in this case, taking into account speeds of rotation of disk drives included in the group of disk drives 107 as well, the disk control hypervisor may allocate the group of disk drives 107 having a disk drive of a high rotation speed, for example, 15,000 rpm among the

groups of disk drives 107 of the same RAID1 configuration to the logical partition 1.

Note that allocation of the host interface unit 100 to logical partitions is performed by the file control hypervisor according to performance required by a host system using the each logical partition. More specifically, in the case in which the performance required by the host system is high, the file control hypervisor allocates a large amount to a logical partition used by the host system, that is, high ability of communication with the host system (communication bandwidth, etc.). On the other hand, in the case in which the performance required by the host system is low or, in particular, there is no request, it is conceivable that the file control hypervisor allocates a small amount to a logical partition used by the host system, that is, low ability of communication with the host system and improve efficiency of the entire NAS.

Moreover, by logically partitioning to use one NAS as in this embodiment, it is possible to perform user authentication for a host system in the NAS independently by each logical partition. Fig. 6 is a diagram showing an example of the user authentication.

In this figure, a user A with an identifier (hereinafter referred to as "ID") "abc" and a user B with an ID "def" use a host system A which uses a logical partition 1, and a user C with an ID "ghi" and a user D with an ID "abc" use a host

system B which uses a logical partition 2. In this case, the user A of the host system A and the user D of the host system B have the same ID "abc". Thus, in order to distinguish the user B and the user D in a conventional NAS, it has been necessary to perform special processing such as giving IDs to host systems or a group of host systems and combining the host ID and IDs of users to distinguish the users.

However, in this embodiment, the host interface unit 100 and the file control processor 101 are logically partitioned, and the each logical partition operates as the host interface units 100 and the file control processors 101 of separate individual virtual NASs. Thus, user authentication is also performed independently for each logical partition. In other words, the user A and the user D having the identical ID "abc" in Fig. 8 are authenticated in the separate logical partitions respectively. Therefore, the user A and the user D are distinguished naturally, and no special processing is required in order to distinguish the users. In other words, as long as logical partitions are different, an identical ID can be given to plural users without performing special processing.

Moreover, resources of the host interface unit 100 and the file control processor 101 allocated to each logical partitions are never used in the other logical partitions. Thus, even if a user of a certain logical partition performs a large quantity of data access, users of the other logical partitions

are never affected by that data access.

Next, a second embodiment will be described. A NAS of this embodiment includes processors of one type and memories of one type, in which the file control processors 101 and the disk control processors 104 of the NAS are integrated and the file control memories 102 and the cache memories 105 of the NAS are integrated.

Fig. 7 is a diagram showing an example of a structure of the second embodiment. In Fig. 7, integrated control processors 901 are processors in which the file control processors 101 and the disk control processors 104 are integrated, and integrated control memories 902 are memories in which the file control memories 102 and the cache memories 105 are integrated.

Compared with the previously-mentioned embodiment (Fig. 1), in this embodiment, the inter-processor communication units 103 become unnecessary and the system structure is simplified. In Fig. 7, both of processing for translating an input/output request by a unit of file from a host system into an input/output request of data and control of input/output of data between the disk interface units 106 and the groups of disk drives 107 are performed by the integrated processors 901. In addition, translation control information and data of the groups of disk drives 107 are stored in the integrated control memories 902. Structures and operations of the other portions in Fig. 7 are

the same as those in Fig. 1.

In Fig. 7, as in Fig. 1, for example, processing is performed independently in logical partitions separated by a logical partition boundary as shown in Fig. 7. When the host interface units 100, the integrated control processors 901, the integrated control memories 902, the disk interface units 106, and the groups of disk drives 107, all which are physical resources for processing allocated to the respective logical partitions, are allocated to each logical partition once, the devices are used solely for processing of the logical partition. In this way, the respective logical partitions operate as virtual NASs independent from each other.

The processing for logical partitioning and allocating the physical resources to the respective logical partitions is actually executed by the integrated control processors 901. The integrated control processors 901 perform control of the previously-mentioned integrated control hypervisor.

Fig. 8 is a diagram showing an example of a setting input screen of a supervising terminal for inputting setting of logical partitioning of a NAS. Such a setting input screen is used in both the first embodiment and the second embodiment. A supervisor or the like inputs setting for logical partitioning as shown in Fig. 8, and contents of the input of setting are notified to the NAS, whereby hypervisor operating in the NAS logically partitions the respective resources of the NAS. More

specifically, contents of the input of setting are stored in a certain storage area provided in the NAS, and the hypervisor performs logical partitioning in accordance with the stored information.

More specifically, such a supervising terminal may be provided in a host system which has connection to the NAS through a network. Or, the supervising terminal may be realized by a console device connected to the NAS by a special line. Such a control device can be realized by an input device such as a keyboard and a display device such as a display.

Information inputted by the supervisor or the like is transferred to the NAS from the host system or the console device using a special purpose protocol or a general purpose protocol. The NAS has an interface for receiving the information (e.g., the host interface units 100 or a special purpose interface).

The setting input screen shown in Fig. 8 will be hereinafter described in detail. The screen includes a field in which the number of partitions of logical partitioning is inputted. The supervisor or the like inputs an intended number of partitions of logical partitioning in this field first (3 in the example of Fig. 8). When the supervisor or the like inputs the number of partitions of logical partitioning, logical partitions corresponding to the number are displayed on the screen for each physical resource, and an initial value of resource allocation to each logical partition is displayed.

Thereafter, the supervisor or the like inputs allocations of the resources such as processors or memories while looking at the screen. In this case, for example, in a part where allocation of the file control processor 101 and the disk control processor 104 is designated, if display is adapted such that ways of allocation of the respective processors to the respective logical partitions can be set in association with each other as shown in the figure (in the figure, both the processors are displayed side by side such that a relation between the processors can be easily seen), it becomes easy to input the setting for allocation control as described above with reference to Fig. 5.

In addition, for example, as shown in Fig. 8, not only knobs for setting amounts of allocation of the file control processor 101 and the disk control processor 104 to the logical partitions individually (icons which can be selected by a pointing device) but also knobs with which the amounts of allocation can be set in association with each other are prepared. Similarly, concerning the file control memory 102 and the cache memory 105, amounts of allocation of the memories to the respective logical partitions are displayed in association with each other, and knobs for individual setting and associated setting are prepared.

In the example of Fig. 8, concerning the inter-processor communication unit 103, the supervisor or the like inputs

percentages of allocation of entire data transfer ability between the file control processor 101 and the disk control processor 104 (the host interface unit 100 and the cache memory 105) to the respective logical partitions.

In addition, concerning the host interface unit 100, the supervisor or the like inputs information on allocation by specifying a connection port of a network for physically connecting the host system rather than according to percentages of the resources. However, a method of inputting the information simply according to the percentages of the resources may be adopted. Moreover, concerning the group of disk drives 107, the supervisor or the like performs allocation of the resources by setting capacities, RAID constitutions, and performance (speed of rotations) of physical disk drives to be allocated to the respective logical partitions.

The above-mentioned methods of input of setting for resource allocation are examples. Other than these methods of input, allocation of resources may be inputted independently as numerical values or may be inputted automatically to some extent. For example, the supervisor or the like inputs characteristics of access which a host system requires of a certain logical partition (random or sequential, an average data length per one transfer, a minimum data transfer rate, a maximum response time, etc.) via a supervising terminal, whereby the supervisor or the like selects a set of parameters

meeting the inputted characteristics with hypervisor, from several sets of parameters prepared in advance, more specifically, from sets of parameters having the characteristics shown in Figs. 2 to 5.

Consequently, for example, in the case in which the supervisor or the like designates sequential access, the supervisor or the like selects a set of parameters corresponding to sequential (parameters corresponding to sequential shown in Figs. 2 and 4) with hypervisor from the sets of parameters prepared in advance and sets the parameters to perform processing of logical partitioning.

More specifically, when the supervisor or the like wishes to create three logical partitions, namely, a logical partition suitable for a large file access, a logical partition suitable for small file access, and a moderate size file access, the supervisor or the like prepares partitioning parameters for a logical partition 1, in which an amount of allocation of a file control processor is smaller than an amount of allocation of a disk control processor, a logical partition 2, in which an amount of allocation of a file control processor is larger than an amount of allocation of a disk control processor, and a logical partition 3, in which amounts of allocation of a file control processor and a disk control processor are comparable, as described in the example of allocation setting of a file control processor and a disk control processor in Fig. 8, in

an area referred to by the hypervisor (in a storage area in which information on allocation of resources to logical partitions described later in Fig. 9 is stored).

In actually performing allocation setting, the supervisor or the like simply designates characteristics which are requested of logical partitions such as the three logical partitions consisting of the logical partition suitable for large file access, the logical partition for small file access, and the logical partition for moderate size file access. Then, the hypervisor automatically selects parameters for allocation corresponding to the designation.

Consequently, the supervisor or the like can easily designate logical partitions having desired performance and characteristics.

In addition, the supervisor or the like is required to make sure such that resources sufficient for the each logical partition allowing it to operate correctly are always allocated by the input of allocation setting. For example, an amount of allocation of the file control processor or the disk control processor cannot be set to zero. At this point, in the automatic setting as described above, it is also possible that a lower limit is set for an amount of allocated resources in advance such that this lower limit is complied with automatically. In addition, in the example of input as shown in Fig. 8, it is also possible that lower limits of amounts of allocation of

the respective resources are defined in advance in the NAS and, in the case in which an amount of allocation of a resource below the lower limit for the resource is inputted, an warning is reported or such input of allocation is checked and refused.

Consequently, the supervisor or the like can set logical partitions safely.

Fig. 9 is an example showing information indicating the allocation of the resources to the logical partitions set by the supervisor or the like as described above. A correspondence relation between the physical resources and the logical partitions as shown in Fig. 9 is created on the basis of information that the hypervisor has received from the supervising terminal. More specifically, the respective parts of hypervisor have information on structures of the physical resources of the NAS, allocate the physical resources to the respective logical partitions on the basis of information inputted by the supervisor or the like and the information on the structure, and create the correspondence relation as shown in Fig. 9. Note that, in Fig. 9, the number of items on the vertical axis increases and decreases according to structures of devices provided in the NAS, and the number of logical partitions on the horizontal axis is changed according to designation by the supervisor or the like.

Then, the information on the correspondence relation as shown in Fig. 9 is stored in an area solely used by hypervisor

of any one or more of the file control memories 102, the cache memories 105, the groups of disk drives 107, and the other storage media as described above. The respective parts of hypervisor determine resources, which the respective logical partitions are caused to use, with reference to the stored information and perform processing for allocation of the resources.

In the storage according to the present invention, when the storage is shared by plural hosts, control for the sharing can be reduced, and mutual interference among the hosts can be eliminated to guarantee performance of data input/output and localize data destruction or failure.

In addition, user authentication independent for each host or each group of hosts can be performed. Further, according to the present invention, it is possible to improve utilization of resources such as processors, memories, and storage media in a system.

Moreover, plural virtual NASs can be provided by one NAS. A degree of freedom of an operating system of a host can be improved. It is possible to independently perform operation and stop or failure processing. Thus, the NAS becomes advantageous in terms of system management, an installation area, power consumption, and so on.

Fig. 10 is an example of a form in which an NAS is connected to a host system (host computer). The NAS according to the present invention can also be used in the form of Fig. 10.

In Fig. 10, four host computers 900a to 900d are connected to an NAS 902 by a network 901. A supervising terminal 903 is connected to the NAS 902 by a special line 904. By using the special line 904, the supervising terminal 903 can be connected to the NAS 902 even if no settings are made in the NAS 902 in advance. When the NAS 902 is used for the first time, since no settings are made concerning a network, it is possible that the supervising terminal 903 is connected in the form as shown in Fig. 10.

On the other hand, Fig. 11 is an example of a form in which the supervising terminal 903 is also connected to the NAS 902 through the network 901 without providing a special line between the supervising terminal 903 and the NAS 902. In this case, settings concerning a network have to be made in the NAS 902 in order to communicate with the supervising terminal 903 through the network. If the network 901 is, for example, an IP network, settings for IP addresses of the NAS 902 itself and the supervising terminal 903 and network masks are necessary.

It is possible that such settings concerning a network for communicating with the supervising terminal 903 are performed by, for example, connecting the supervising terminal 903 to the NAS 902 once through a special line in the form as shown in Fig. 10. When the settings for the network are completed, the connection by the special line between the supervising terminal 903 and the NAS 902 can be cancelled and removed, and

the supervising terminal 903 can be connected to the network 901 to change the form to the form of Fig. 11, whereby the NAS 902 can be supervised from the supervising terminal 903 through the network.

In addition, in Fig. 11, another method of performing the settings for the network for communicating with the supervising terminal 903 is to install a very small console unit for only performing the settings for the network in the NAS 902. Fig. 12 is an example of such a console unit. If the console unit as shown in Fig. 12 is provided on an appropriate surface of a housing of the NAS 902, the network settings for communicating with the supervising terminal 903 through the network 901 in the form as shown in Fig. 11 can be performed. If the supervising terminal 903 can be connected once through the network 901, supervising work for the NAS 902 after that can be performed through the supervising terminal 903.

In the NAS according to the present invention, plural virtual NASs operate on one physical NAS. For identifying the respective virtual NASs, settings are made in principle in network connection ports (host channels) of host interface units, which are allocated to the respective virtual NASs (logical partitions), such that the virtual NASs are identified by the network individually. For example, in the case in which the virtual NASs make connection through an IP network, different IP addresses have to be assigned to the respective host channels.

Such identification settings for the network (assigning of IP addresses) can be performed by making connections between the supervising terminal and each of the virtual NASs operating in each logical partition. When the supervising terminal is connected to the physical NAS, for example, by a special line as shown in Fig. 10, it is possible to provide a special switch on the supervising terminal side or the physical NAS side for switching the virtual NAS to be connected to the supervising terminal. For example, it is possible that such a switch is adapted such that currently operating virtual NASs are switched to be connected to the supervising terminal in turn every time the switch is pressed. On the supervising terminal side, it is also possible that such a special switch is substituted with some special sequence of normal key switches on the supervising terminal.

In addition, in the case in which the supervising terminal is connected through the network as shown in Fig. 11, identification settings for supervising (if the network is an IP network, assignments of IP addresses) different from identification settings of the network given to the host channels are performed in the NAS. In this case, first, the supervising terminal makes connection to the NAS using an IP address for supervision. Next, the special switch is provided such that currently operating virtual NASs are switched in turn every time the switch is pressed. Alternatively, it is also possible

to prepare IP addresses for supervision by the number of virtual NASs (the number of logical partitions). In that case, it is unnecessary to provide a special switch, but it is necessary to prepare addresses of the network for supervision by the number of virtual NASs.